ESD-TR-81-87

Semiannual Technical Summary

Packet Speech Systems Technology

31 March 1981

Prepared for the Defense Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-80-C-0002 by

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Defense Advanced Research Projects Agency under Air Force Contract F19628-80-C-0002 (ARPA Order 3673).

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Raymond L. Loiselle, Lt.Col., USAF
Chief, ESD Lincoln Laboratory Project Office

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LINCOLN LABORATORY

(6) **PACKET SPEECH SYSTEMS TECHNOLOGY**

(9) **SEMIANNUAL TECHNICAL SUMMARY REPORT TO THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY**

(10) Clifford J. Weinstein
Peter E. Blankenship

1 OCT [redacted]80 – 31 MAR [redacted]981

ISSUED 10 JULY 1981

(12) 26

(11) 31 Mar 81

(15) F19628-80-C-0002, ARPA Order-3673

(16) 1D10, 1T10

(18) ESD

A

(19) TR-81-87

LEXINGTON MASSACHUSETTS

ABSTRACT

This report describes work performed on the Packet Speech Systems Technology Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 October 1980 through 31 March 1981.

CONTENTS

	Abstract	iii
	Introduction and Summary	vii
I.	DIGITAL CHANNEL VOCODER	1
II.	EMBEDDED CVSD-BASED WAVEFORM CODER	7
III.	PACKET VOICE TERMINAL AND LOCAL ACCESS AREA	9
	A. Introduction and Summary	9
	B. Packet Voice Terminal Hardware	9
	C. Packet Voice Terminal Software	11
	D. LEXNET Traffic Emulation and Measurement	11
	E. Replication and Technology Transfer	12
IV.	MINICONCENTRATOR	13
	A. Introduction and Summary	13
	B. Miniconcentrator Hardware	13
	C. Miniconcentrator Software	14
	D. Multiport Memory Design	16
V.	EXPERIMENT DEFINITION AND PLANNING	17

INTRODUCTION AND SUMMARY

The long-range objectives of the Packet Speech Systems Technology Program are to develop and demonstrate techniques for efficient digital speech communication on networks suitable for both voice and data, and to investigate and develop techniques for integrated voice and data communication in packetized networks, including wideband common-user satellite links. Specific areas of concern are: the concentration of statistically fluctuating volumes of voice traffic, the adaptation of communication strategies to varying conditions of network links and traffic volume, and the interconnection of wideband satellite networks to terrestrial systems.

Previous efforts in this area have led to new vocoder structures for improved narrowband voice performance and multiple-rate transmission, and to demonstrations of conversational speech and conferencing on the ARPANET and the Atlantic Packet Satellite Network.

The current program has two major thrusts: i.e., the development and refinement of practical low-cost, robust, narrowband, and variable-rate speech algorithms and voice terminal structures; and the establishment of an experimental wideband satellite network to serve as a unique facility for the realistic investigation of voice/data networking strategies.

This report covers work in the following areas: digital channel vocoder development; embedded CVSD-based speech waveform encoder design and implementation; development and experimental tests of modular packet voice terminals (PVTs) and local access area (LEXNET) facilities; development of a miniconcentrator facility to mediate the flow of traffic from the LEXNET onto the wideband satellite network, and execution of packet speech experiments using this facility; and definition and planning of, and participation in, experiments on the wideband integrated voice/data network.

The hardware design of a single-board Belgard-type digital channel vocoder with Gold pitch detection has been completed. The design is built around soon-to-be-available commercial general purpose digital signal processing chips. Key microcode routines for these chips have been coded and timed, and critical software development tools have been implemented. A compact embedded data stream speech waveform coder has been designed to support rate-adaptive speech flow control experiments with the PVT. The coder supports embedded rates of 16, 32, 48, and 64 kbps, and is compatible with standard CVSD systems at 16 kbps.

Thorough integration and test of second-generation PVTs have been carried out, and both LPC and PCM capability have been demonstrated with Network Voice Protocol (NVP-2) and stream (ST) protocol software operating in PVTs. Replication of PVT and LEXNET hardware has proceeded, and a LEXNET with two PVTs has been installed and successfully operated at Information Sciences Institute (ISI). The miniconcentrator became operational as an internet gateway during this reporting period and was successfully demonstrated in January as a LEXNET-to-ARPANET gateway for LPC speech between a LEXNET PVT at Lincoln and an

ARPANET speech host at ISI. A detailed system architecture has been developed for an intelligent multiport buffer memory to handle address manipulation and packet throughput tasks for the miniconcentrator. Experiment planning and coordination activities, including publication of the second annual experiment plan supplement, have continued. Three of the wideband satellite (WB SATNET) network nodes – Lincoln, ISI, and Defense Communications Engineering Center – are now fully equipped with demand assignment processors, burst modems, and earth stations; integration and test of the WB SATNET is in progress.

PACKET SPEECH SYSTEMS TECHNOLOGY

I. DIGITAL CHANNEL VOCODER

The hardware design has been completed for a Belgard-type, 10-kHz, 19-channel filter bank vocoder which interfaces to the Packet Voice Terminal (PVT) protocol processor. It is implemented with 26 commercially available or soon-to-be-available integrated circuits, dissipates an estimated 11.5 W of power, and occupies about 70 percent of a 7-in. × 7-in. Augat universal wirewrap board. A multiprocessor, distributed signal processing architecture is used based on six Nippon Electric Company (NEC) μPD 7720 signal processing chips and an INTEL 8085-based microcomputer for system control and protocol processor communications. Innovative use of CODEC-with-filters LSI integrated circuits results in a compact analog subsystem which in a vocoder application exhibits quality equivalent to the LDSP laboratory-reference signal conditioner. A feature of the design is that by reprogramming the NEC processors and control microcomputer EPROMs, and changing the sampling rate crystal, the identical hardware will implement an 8-kHz LPC-10 algorithm. Furthermore, preliminary benchmarking shows that NEC ROM resources are sufficient to simultaneously accommodate both LPC and Belgard firmware, indicating the possibility of a user-selectable vocoder. Also, NEC μPD 7720 software development aids have been written including a cross-assembler and a non-real-time simulator/debugger running under the UNIX operating system. These programs have been used to debug the code for the three NEC chips implementing the Belgard spectrum analysis. Debugging of a one-NEC-chip Gold pitch detector is in progress.

The channel vocoder architecture is shown in Fig. 1. The analog input speech is processed through the analog-to-digital conversion subsystem which produces a serial data stream for input to a 4-chip analyzer array. The analyzer signal processing task is distributed such that one NEC chip computes the weights for the lower seven spectral channels, while two NEC signal processing chips compute six more spectral channel weights each. A fourth NEC chip implements the Gold pitch detector providing the frame-wise pitch estimate and voicing decision. Each analysis frame, the resulting coded 19 spectral channel weights and 2 pitch detector values are transferred over an 8-bit (INTEL 8080) microcomputer bus under the control of an INTEL 4-chip 8085 microcomputer to the protocol processor. In a similar fashion, the control processor receives coded synthesis parameters each frame from the protocol processor and transmits them to an NEC-chip synthesizer array. One μPD 7720 synthesizes the lower ten spectral channels of the output speech, while a second chip synthesizes the remaining upper nine spectral channels. By making use of the NEC chip's serial input and output ports, the upper and lower spectral portions are digitally summed and transferred to the digital-to-analog converter.

In recent hardware vocoder implementations, analog subsystems have required a dozen or more packages, many of them performing low-level or SSI (Small Scale Integration) functions. In this design, where the digital signal processing task itself requires only six 28-pin DIPs, such an analog/digital conversion system would account for an undue proportion of circuit board "real estate." A CODEC-with-filters integrated circuit (e.g., AMI S3505, Hitachi HD44212) was viewed as offering a possible LSI (Large Scale Integration) solution to this problem. In one 24-pin package, the Hitachi and AMI chips include a complete analog-digital conversion system (anti-aliasing filter and ADC) and digital-analog system (DAC and post-sampling filter). The CODEC-with-filters chips, though, are designed to meet telecommunications industry standards for PCM

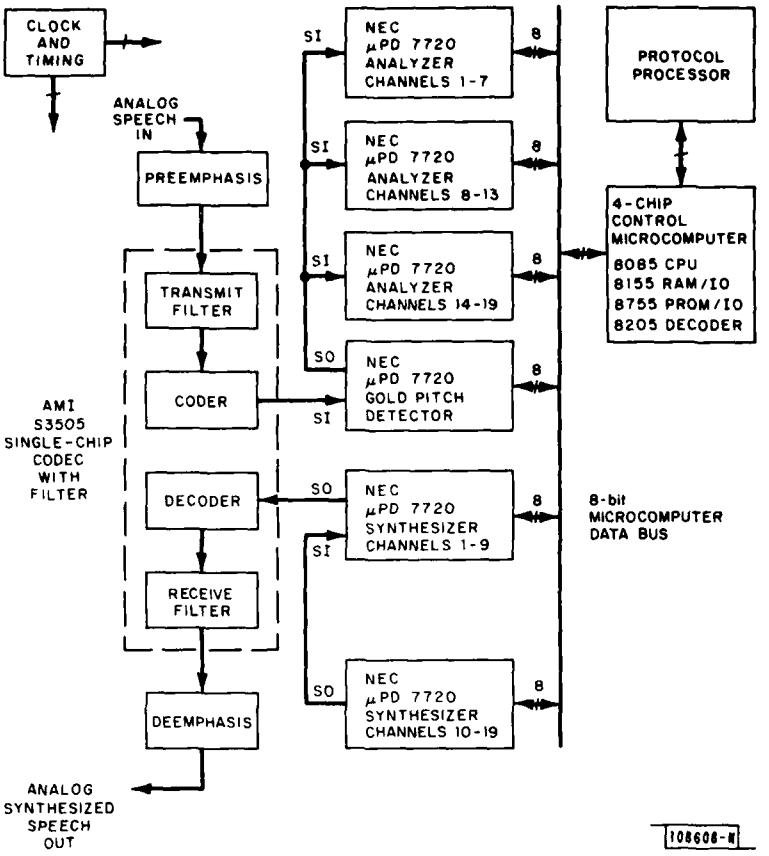

Fig. 1. NEC μPD 7720-based Belgard-type channel vocoder.

communications rather than to serve as a vocoder analog interface. Three specifications were of particular concern:

(1) The coder (A-D converter) and decoder (D-A converter) obey an 8-bit quasi-log (mu-255) law rather than a standard 12-bit linear characteristic.

(2) The transmit (anti-aliasing) filter has a 60-Hz rejection filter resulting in a low-frequency rolloff of 20 dB between 60 and 200 Hz which could potentially degrade the pitch detector's performance.

(3) The transmit and receive (post sampling) filters provide only 14 dB of rejection at the folding frequency (4 kHz) which could result in aliasing noise detrimental to vocoder performance.

To evaluate the actual performance of these chips, real-time LDSP-based "A-B" comparisons were made between a Lincoln LPC-10 using the Hitachi and AMI-based analog subsystems and an LPC-10 equipped with the standard signal conditioner as a reference. Informal listening tests by trained listeners revealed no detrimental effect on vocoder quality when either of the CODEC-with-filters-based analog subsystems were substituted for the reference. Significantly,

the pitch detector performance did not seem to be degraded by compromising the low-frequency information. In fact, the CODEC systems were noted to produce greater "clarity" in the synthetic speech. This is most likely due to the sin x/x correction, implemented in the CODEC smoothing filters but not in the reference, which results in a restoration of some high-frequency information. Finally, DRT tests were found to support the results of the listening sessions: a three-male speaker DRT score of 86.8 (standard error = 0.59) was obtained for the HD44212/LPC-10 combination while the reference signal conditioner/LPC-10 combination scored a nearly equivalent 88.1 (standard error = 0.79).

Due to uncertainty in future availability of the Hitachi CODEC-with-filters chip, the AMI version was chosen for the actual channel vocoder hardware. To minimize processing overhead in this design, the 8-bit mu-law-coded input speech from the AMI S3505 coder (A-D converter) is decoded once by the pitch detector chip and then, utilizing the NEC serial data ports, is distributed as a 16-bit linear value to the remaining three analyzer chips. Conversely, the 16-bit linear synthesized output speech is recoded in an 8-bit mu-law format by one of the synthesizer chips and transmitted to the AMI S3505 decoder (D-A converter).

An actual-size layout of the channel vocoder on an Augat universal wirewrap board indicating major functional partitioning is shown in Fig. 2. The corresponding package count and estimated power dissipation is given below:

BELGARD PACKAGE COUNT AND POWER DISSIPATION

Device	Quantity	Power Dissipation
40-pin INTEL 8085 CPU	1	1.5
40-pin INTEL 8155 RAM/IO	1	1.5
40-pin INTEL 8755 PROM/IO	1	1.5
28-pin NEC μPD 7720 Signal Processing Chip	6	5.4
24-pin AMI S3505 CODEC with Filter	1	0.1
14- and 16-pin Miscellaneous (clocks, timing, etc.)	16	1.5
Total	26 Packages	11.5 W

An LPC-10 algorithm can be implemented by the same architecture as the channel vocoder, but will require only one NEC chip for spectral analysis and one NEC chip for synthesis. An INTEL 8051 single-chip microcomputer could ultimately replace the 4-chip 8085-based control microcomputer, but will not be considered until 8051 software development aids become available. In fact, by reprogramming NEC and control microcomputer EPROMs, and switching sampling rate crystals, the Belgard hardware will implement the LPC algorithm. Furthermore, initial code sizing indicates that NEC program and data ROMs are large enough to encompass both vocoder algorithms concurrently, resulting in a run-time selectable vocoder.

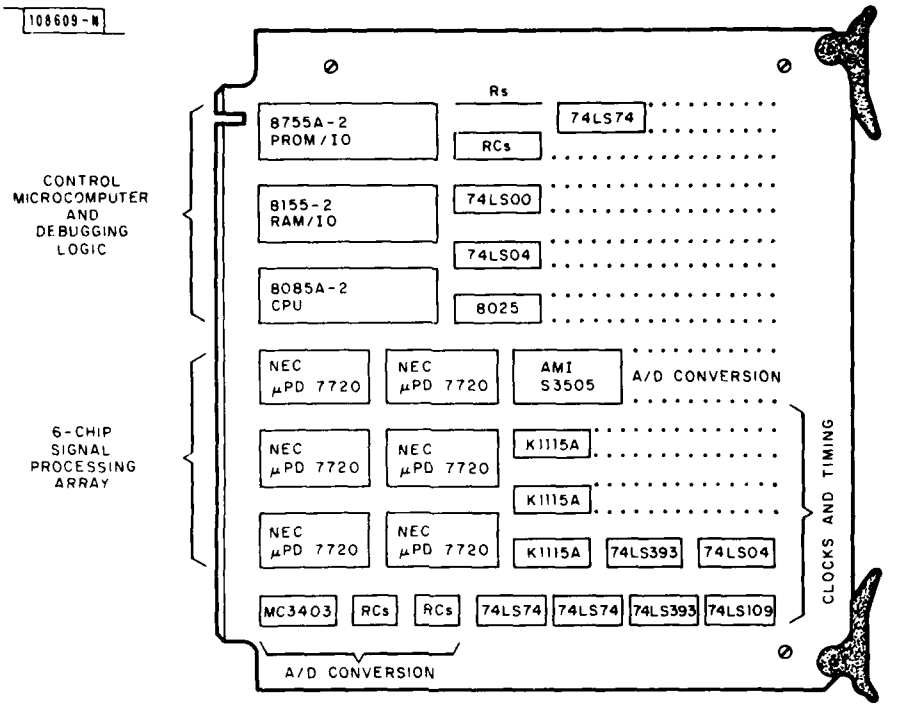

Fig. 2. Single-board (7 in. × 7 in.) Belgard vocoder.

A cross-assembler (NAS) and a non-real-time simulator/debugger (NECSIM) for the 7720 have been written to run on the Group's PDP-11/45 and 11/70 under the UNIX Version VII operating system since no NEC software aids are available outside the Laboratory. The cross-assembler is written in Macro-11 and was converted from the LDSP assembler while NECSIM was written in the C-programming language. Code development with this facility begins with syntactic debugging of NEC assembly language code using NAS. When an error-free assembly is achieved, the resulting binary object code may be used to initialize program and data memory for an NEC program simulation using NECSIM. A typical simulation proceeds in two modes: "single-step" and "continuous" execution. In single-step mode, instructions are executed one at a time returning to "command level" and displaying machine status information between each one. In command level, the user can examine and alter any of the simulated machine registers or memory. Continuous-mode simulation executes program memory code until a user-specified program counter breakpoint is encountered at which time the simulation returns to command level. From command level, execution may be reinitiated in either single-step or continuous mode. A powerful feature of the simulator is its use of the UNIX file system to simulate input/output: NEC read and write instructions to the NEC serial and parallel data ports result in sequential writes and reads, respectively, to and from user-specified data files in the operating system. Other features of the simulator include periodic (user-specified) inputs, outputs, and simulated interrupts as well as locked-out interrupt and input and output register overwriting warnings (useful in debugging of real-time-constrained foreground-background structure programs); an execution timer; and subroutine stack underflow and overflow warnings.

TABLE I

BELGARD VOCODER FOREGROUND PROGRAM EXECUTION TIMES (330-μs NEC Cycle)

	Percentage of 100-μs Sampling Interval
Gold Pitch Detector	
Nominal tasks	30
Additional tasks when peak is detected (worst case)	+111
Total worst case foreground	141
Analysis (7 of 19 filter bank channels)	
Nominal tasks	72
Additional tasks at downsampling and frame boundaries	+49
Total worst case analysis foreground	121
Synthesis (10 of 19 filter bank channels)	
Nominal tasks	62
Additional tasks at downsampling and frame boundaries	+67
Total worst case synthesis foreground	129

The NAS and NECSIM programs are presently being used to implement the Belgard vocoder tasks. The Gold pitch detector, which is also used in the LPC vocoder, is implemented with a "background-foreground" program structure. The foreground processing task, which must be completed each sampling interval, is comprised of a nominal data-independent computation executed each sample and additional tasks executed when a peak is detected in the filtered input speech waveform. Although in the worst case, the Belgard foreground program execution time (Table I) actually overruns one sampling interval, NEC chip input buffering relaxes the real-time constraint by allowing the processing load to be averaged over two sampling intervals. Furthermore, since the occurrence of two consecutive peaks is not possible in the Gold algorithm, real time is maintained if the nominal processing load and the worst case processing load combined is less than two sampling intervals which at a 10-kHz sampling frequency (Belgard) is achieved:

Nominal load	29
Worst case load	+131 μs
	160 < 200 μs = 2 sampling intervals

Based on known LDVT (Lincoln Digital Voice Terminal) execution times and similar instruction efficiencies of the NEC and LDVT architectures, the Gold pitch detector background tasks are estimated to require an average of 16.5 μs per sample for a Belgard 200-sample frame which in conjunction with the foreground tasks described above can be accommodated in real time.

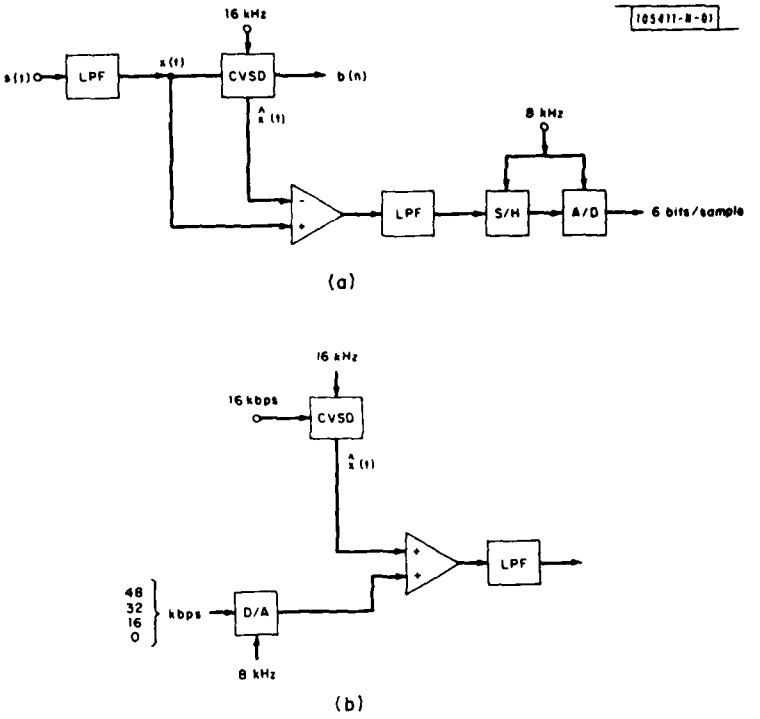

Fig. 3. CVSD-based embedded waveform coder.

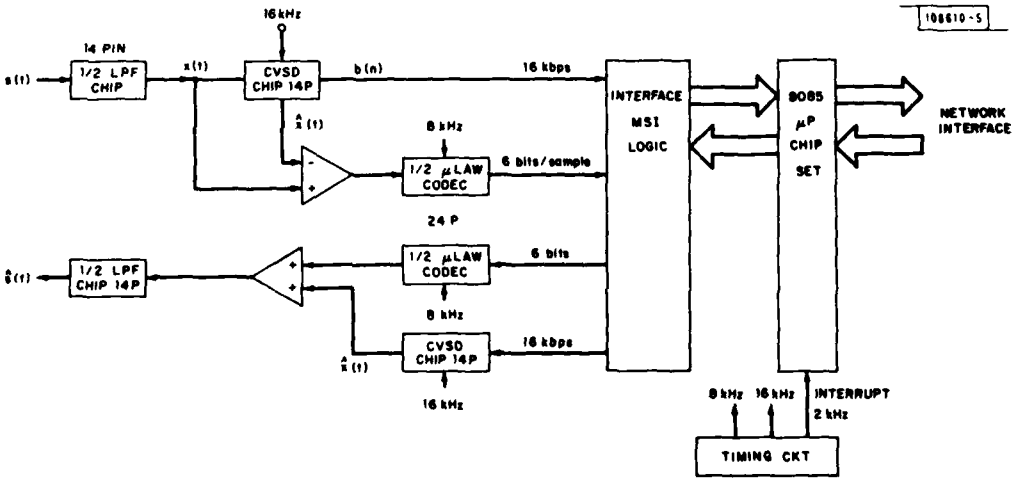

Fig. 4. Hardware realization of embedded CVSD waveform encoder.

The Belgard spectral analysis and synthesis foreground program execution times are shown in Table I for processing loads of seven analysis and ten synthesis filter bank channels (corresponding to the processing distribution in the Belgard design). A sample rate processing load averaging strategy similar to that of the pitch detector is used here at downsampling and frame boundaries where foreground computations may overrun one sampling interval. The Belgard analyzer background computation requires an average of 2 μs/sample. The synthesizer background load is estimated to be a similar value.

In summary, hardware design of a single 7-in. × 7-in. board Belgard vocoder implementation with Gold pitch detection has been completed. The resulting design, utilizing a distributed signal processing architecture, based on six NEC μPD 7720's and an INTEL 8085 microprocessor, is flexible enough to be reprogrammed to implement an LPC-10 algorithm. A compact analog subsystem is achieved with the use of the AMI S3505 CODEC-with-filters chip. A NEC signal processing chip software development facility has been created consisting of an assembler and non-real-time simulator/debugger which has been used to debug the Belgard spectral analyzer code and is presently being used to code a single-NEC-chip Gold pitch algorithm used in both LPC and Belgard analyzers.

II. EMBEDDED CVSD-BASED WAVEFORM CODER

A compact embedded data stream speech waveform coder has been designed to support rate-adaptive speech flow control experiments with the Packet Voice Terminal. This embedded coder provides a speech data stream of 64 kbps in which is embedded speech data streams at 48, 32, and 16 kbps. At the lowest embedded rate of 16 kbps, the encoder produces a data stream strictly compatible with present Defense Communication System 16-kbps CVSD (Continuously Variable Slope Delta Modulation) encoders. The embedded output data structure lends itself to various adaptive flow control strategies used in packet networks. Since the encoded output stream can be transmitted in packets of descending priority starting with the lowest rate data, it is possible to discard lower priority packets when the network traffic will not allow high-rate transmission. The receiver will produce speech output with quality which increases with its received data rate. As shown in Fig. 3, the basic device consists of a CVSD encoder at 16 kHz with filtered input x(t), data stream b(n), and integrated loop signal $\hat{x}(t)$. The backbone encoder is augmented with the open-loop encoding of the filtered difference signal $\Delta = x(t) - \hat{x}(t)$ at an 8-kHz sample rate. This extra information serves as an additive correction to the CVSD approximation to x(t). At the receiver-decoder, the CVSD data stream, b(n), is used to recreate the output signal $\hat{x}(t)$. If no more data is available (i.e., the receiver rate is 16 kbps), the final output is simply the CVSD signal $\hat{x}(t)$ low pass filtered to produce $\hat{s}(t)$.

The 8-kHz samples of the continuous Δ signal must be encoded with some quantizer law. There is not a simple relationship between quantization noise introduced by the Δ samples and quantization in the reconstructed samples using both the CVSD and Δ outputs. Experiments were carried out to select an appropriate quantizer characteristic for the Δ samples. A comparison of fixed linear, fixed logarithmic, and adaptive quantization (based on the CVSD slope waveform) was carried out. The result was that the fixed log encoder and the adaptive scheme showed equivalent performance, both superior to the linear scheme. The fixed log quantizer has been chosen because of its relative simplicity.

Figure 4 presents the hardware realization of the E (embedded) CVSD coder based on currently available PCM, CVSD, and 8085 microprocessor chips. The input and output low-pass

filtering is performed by a single chip containing a pair of switched-capacitor filters designed for commercial PCM applications. Two separate CVSD chips are used, each of which implements the CVSD encoding loop including the two internal smoothing filters. In one case, the chip is used for encoding in the transmit section. The second chip is used in the receive section for decoding. The low-pass filtering of the Δ error signal and then sampling and quantizing to eight bits is performed by one half of a commercial mu-255-law PCM coder chip. Only the more significant six bits are used. At the receive section, digital-to-analog conversion of the Δ signal is performed by the rest of the coder chip. Timing functions and interface logic are implemented with MSI TTL packages. The realization consists of 3 chips from the 8085 μP family, 2 CVSD chips, 1 PCM coder chip, 1 switched-capacitor filter chip, and about 30 MSI TTL chips mounted on a PVT-compatible 7-in. × 7-in. plug-in board. The board will dissipate about 8 W and will plug into the external vocoder slot in the Packet Voice Terminal.

Figure 5 presents the embedded, priority-ordered, packet output from the ECVSD. Each 20-ms interval (chosen to be compatible with frame-oriented vocoder devices), four separate groups of data words are output from the ECVSD to the network terminal. The highest priority packet consists of data bytes composed of 16-kbps CVSD bits ($b^1 \ldots b^8$) as shown at the top of Fig. 5. In the case shown, 40 of these bytes will represent a 20-ms speech waveform at sample

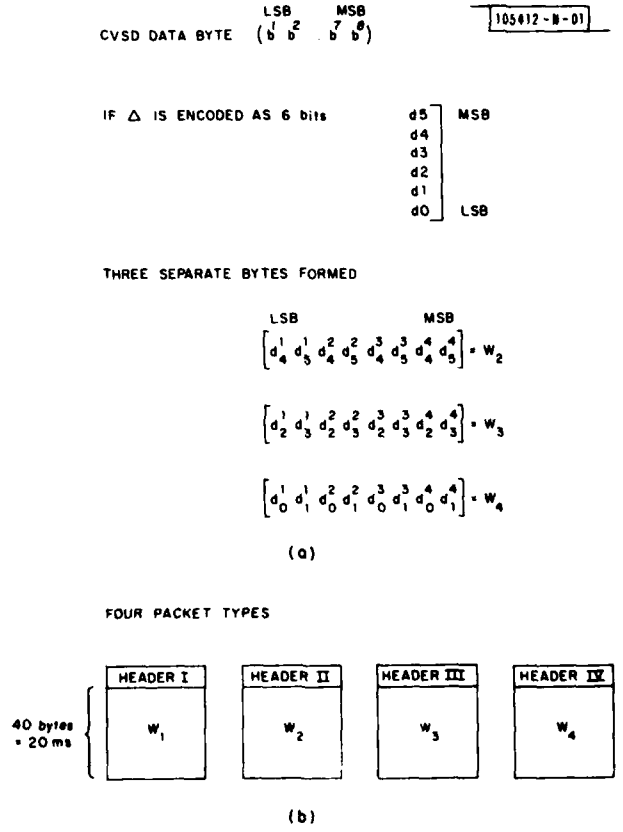

Fig. 5. Output formats for ECVSD encoder.

spacing of 62.5 μs. The remaining three lower priority packets (II, III, and IV at the bottom of Fig. 5) are composed of bytes shown as W_2, W_3, and W_4. They divide the six-bit-encoded Δ signal into three separate parts; d^5d^4, d^3d^2, d^1d^0. Each part is accumulated over four samples to form the bytes shown.

At the receiver, speech is synthesized dependent on the received packets. If only type I and II packets are received, then the output speech will be a 32-kbps waveform consisting of a 16-kbps CVSD waveform enhanced with a two-bit (d^5d^4) Δ correction. If all four packet types are received, the output represents a 64-kbps enhanced CVSD waveform.

Microcode for the 8085 chip set has been blocked out. The code must format the CVSD and error bits into the four packet types and undo the process for the receiver. Silence detection will be done in the 8085 software as well.

A description of the Embedded CVSD coder was presented at the ICASSP meeting in Atlanta and has been published in the ICASSP proceedings for April 1981.

III. PACKET VOICE TERMINAL AND LOCAL ACCESS AREA

A. Introduction and Summary

The previous Semiannual Technical Summary describes the design and initial implementation of second-generation Packet Voice Terminal (PVT) units and the development of voice protocol software for these PVTs. During the past six months, thorough integration and test of the PVT units and protocol software have been carried out. Both LPC and PCM capability have been demonstrated with Network Voice Protocol (NVP-2) and stream (ST) protocol software operating in PVTs. A pair of PVTs and associated Lincoln Experimental Access Area Network (LEXNET) interfaces have been installed and successfully operated at ISI. Replication of PVT and LEXNET hardware has proceeded, with a total of five PVT units and three LEXNET-to-concentrator interface (LCI) units now constructed. In preparation for technology transfer of this hardware, all wirelists have been automated, and detailed schematics and mechanical design layouts are being drawn.

B. Packet Voice Terminal Hardware

A photograph of a complete second-generation PVT unit is shown in Fig. 6. The five occupied card slots in the PVT include the modem, LEXNET buffer control unit, protocol processor, memory extension, and PCM/vocoder interface unit. The sixth card slot is wired to house a single-card vocoder such as the channel vocoder or embedded waveform coder described in earlier sections. Interface to an external vocoder is achieved through the vocoder interface card. The special telephone instrument includes an 8085 microprocessor complex which controls dialing and signalling and communicates with the PVT protocol processor via an RS-232 link.

A number of hardware problems have been encountered and overcome during the debugging of the RAM-resident programs in the protocol processor. One set of problems involved memory faults in the INTEL 8185 RAM program storage. These have been overcome by: (1) discovering a single-bit failure problem, which occurred only after the memory sat idle for 20 to 30 s, in 7 out of a lot of 100 8185 chips; and (2) dropping the clock rate on the INTEL 8085 CPU from 10 to 8 MHz. In the first case, the faulty chips have been removed. In the second case, the lowered clock rate has eliminated failures due to marginal operation of some of the chips at 10 MHz, and the 8-MHz rate is still more than adequate to handle the worst case processing

Fig. 6. Packet Voice Terminal unit.

load which results from 64-kbps PCM. A second set of problems involving the bidirectional bus coupler used to reduce the loading of the 8085 bus by the memory extension card has been discovered and resolved. On certain of these bus coupler chips, the switching transients were fast enough to induce signals on adjacent control lines. The problem was cured by adding ferrite beads to limit the rate of change of the data signals. This change has been installed on all memory extension cards, and no further problems have been observed.

Several changes have been made in the mechanical design of the rack for the voice terminals. A larger backplane is being used, which can accommodate nine cards instead of the present six cards. The power supply and front panels are now designed as independent modules which plug into the chassis. This should speed fabrication and simplify servicing or modifications should the need arise.

Firmware has been developed and tested for the 8085 processor in the telephone handset. This firmware interprets dialing keys and the off-hook condition, and controls the dial tone, ringing tone, busy signal, and bell. Initial integration and test of the handset was carried out with a PROM-resident test program in the protocol processor. Later, dialing and signalling software was incorporated into the RAM-resident NVP-2 code.

Several revisions were made in the interface between the protocol processor and the PCM card to meet the requirements of the ISI Switched Telephone Network (STN) card which fits into the same slot as the PCM card. These changes have no impact on the normal operation of the PVT.

During January the full LEXNET using the NVP-2/ST protocol, integrated with a LEXNET-to-ARPANET gateway implemented in the PDP-11/45, was used to transmit 2400-bps LPC speech to ISI via the ARPANET. At Lincoln, the speech was generated in an LDSP interfaced to the external speech processor part of the LEXNET Packet Voice Terminal. At ISI, the same algorithm was implemented in an AP-120B.

A LEXNET consisting of two Packet Voice Terminals and a concentrator interface was shipped to ISI on 25 February and was installed the week of 2 March. The NVP code was downloaded into the PVT from the ISI PDP-11/45 and demonstrated using the internal PCM capability of the terminal. In addition, the PVT was interfaced with the ISI STN interface card, which replaces the PCM card. Successful speech transmission was demonstrated over a telephone line to the STN card and then over the LEXNET to a PVT using its internal PCM card.

C. Packet Voice Terminal Software

Software for the PVT to implement point-to-point conversations has been completed. The program follows the NVP-2 and ST protocols and can set up and take down a connection between itself and another terminal which follows the same protocols. NVP-2 and ST packet-handling software have been implemented in the protocol processor for both 64-kbps PCM and 2400-bps LPC. The programs are written in a combination of "C" (for the protocol control functions) and "A-natural" (for the real-time speech and I/O handling) and are downloaded from the PDP-11/70 into the PVT RAM via a TTY link to the trap control unit. The programs have been exercised for PVT-to-PVT communication (both LPC and PCM), both through the LEXNET only and looped through the concentrator. In addition, LPC communication was demonstrated in January between a LEXNET PVT and an ISI speech host on the ARPANET.

The PVT user provides control via the touch-tone pad on his special user instrument. The protocol processor communicates with the user instrument via an RS-232 port. Software in the protocol processor has been implemented to set up NVP-2 connections based on user button pushes. Dial tones, ringing signals, busy signals, etc., are provided under program control, causing the PVT to seem like a regular telephone. The PVT can call another PVT on the LEXNET, or it can go through the gateway to a terminal on another net. Each terminal is currently assigned a three-digit phone number. To route a call through the gateway, the caller dials 9 followed by the desired phone number. A more complete addressing plan is currently under discussion by a number of organizations involved in the wideband network experiment. When decisions on network addressing are reached, the PVT software will be updated in accordance with the selected addressing scheme.

D. LEXNET Traffic Emulation and Measurement

The PVT software has been modified to produce a Traffic Emulation Module (TEM). The TEM is currently controlled by the PVT phones, and an experiment can be run between any pair of PVTs, each downloaded with the TEM code.

An open-ended set of experiment definitions is stored in the PVT memory. To run an experiment, the phone on the "boss" TEM is used to dial three numbers. First the number of the experiment which the "boss" TEM should perform, then the number of the experiment which the "slave" TEM should perform, and finally the phone number of the "slave" TEM. The "boss" then initiates a call to the "slave." During the NVP-2 stage of the protocol, the "boss" tells the "slave" the number of its experiment. The "slave" declares itself ready without ringing the phone.

This mechanism is now working. However, experiment control is presently limited by the fact that the only clock available for timing purposes is the 20-ms interrupt provided by the PCM card. An experiment is defined as the group of messages which should be sent during each 20-ms interval, the number of these groups which should be sent, the inter-message time between messages in a group, and the procedure to follow if a message group is not completed in one 20-ms interval.

This allows any mixture of IP and ST messages, each of arbitrary length, to be sent. As a temporary way to time inter-message gaps, the program counts the number of times it has gone around its main loop. When a message group has not been completed during one interval, the TEM can attempt to maintain the transmission rate by immediately beginning the next group or it can allow time to slip by waiting for the start of the next 20-ms interval.

Each TEM counts the number of messages of each type it receives and the position number within the group of each received message. When sending, it records the number of times a group was not sent within the 20-ms interval. It also tries to determine how often it could not send because the buffer control card was not ready. Beginning and ending time stamps are recorded.

Experiments are just beginning which should provide information on the performance characteristics of the PVTs, the LEXNET, and the Gateway. Several improvements are planned. The protocol processor card has a timer available which will be wired to a clock. This will allow fairly accurate timing of small intervals. At the end of an experiment, the "boss" TEM will send a special "Give Me Your Data" request to the "slave" TEM which will then send back its statistics. A PDP-11/70 terminal will replace the phone as the method of running the "boss" TEM. The current downloading capability will be expanded so that all input to the TEM will be typed in, and the "boss" TEM will use this route to send back statistics. This will provide a record of the experiment since the downloading program puts everything it receives into a file as well as displaying it on the scope.

E. Replication and Technology Transfer

A careful study of production requirements indicated a need for an intensive technician effort as well as the use of semiautomatic wirewrap facilities to expedite manufacture of the PVTs, concentrator interfaces, and trap control cards required for the multi-user packet speech experiments. Due to the critical schedule of delivery commitments, it was decided to circumvent the full-scale use of newly initiated in-house support facilities for documentation and replication. A compromise solution was employed whereby hand wirelisting of wirewrap cards was initiated using formats compatible with the semiautomatic facilities, and all wirewrap boards will be replicated using these facilities. A parallel effort was initiated to create a data base for full-scale documentation, placing, and routing of wirewrap panel designs, as an aid to long-term technology transfer of all items.

Wirelists of all wirewrap cards and the PVT backplane have now been completed and entered into the semiautomatic facility. At least one card of each type has been wired from the list using the semiautomatic wirewrap machine. Automatically wired versions of all cards, except for the PCM card, have been debugged.

The mechanical components for all terminals have been fabricated. A full set of draft schematic and mechanical drawings for the PVT has been submitted to our Engineering Division for the preparation of high-quality drawings suitable for technology transfer.

The construction goals for FY 81 include seven PVTs and two LEXNET-Concentrator Interfaces (LCIs), in addition to the four PVTs and one LCI built during FY 80. Currently, one new PVT and one new LCI have been completed, and one more of each is being assembled. A large number of parts for the other units have been wired. A current summary of PVT/LEXNET construction status is given in Table II. The left-most column refers to the total number of each component required by the end of FY 81. The other two columns indicate how many of each component were wired and tested by 31 March 1981.

TABLE II
SUMMARY OF LEXNET CONSTRUCTION (31 March 1981)

	Total	Wired	Tested
Cable Tap	14	8	4
Modem Card	14	14	9
Buffer Control	14	10	10
Protocol Processor	11	7	7
Memory Extension	11	7	7
PCM Card	11	5	5
Concentrator Interface Card	3	3	2
Telephones	11	6	5

IV. MINICONCENTRATOR

A. Introduction and Summary

The miniconcentrator became operational as an internet gateway during this reporting period, and was successfully demonstrated in January as a LEXNET-to-ARPANET gateway for LPC speech between a LEXNET PVT at Lincoln and an ARPANET speech host at ISI. Miniconcentrator software was successfully exported to a PDP-11/45 at ISI and was utilized there in March for a PCM speech loop test from one PVT on the ISI LEXNET, through the gateway, to the second PVT at ISI. Although the LEXNET-to-concentrator interface has operated satisfactorily in these tests, we are still encountering difficulties in achieving fully satisfactory bit-rate performance through the UMC-Z80 SIO chip. The UMC-Z80 software on the PSAT side has been completely revised to be in the same form as the LEXNET UMC-Z80 software and to be compatible with the PDP-11 I/O developed for the LEXNET. System-level checkout of the miniconcentrator in LEXNET-to-PSAT gateway mode is under way. A detailed system architecture for an intelligent multiport buffer memory (MPM) has been developed. This memory will unburden the Z80s on each net from address manipulation tasks as well as avoid the need for the data portions of packets to pass through PDP-11 memory. The following sections update progress and status on miniconcentrator hardware and software, and on the MPM.

B. Miniconcentrator Hardware

The initial demonstrations of the miniconcentrator took place using the PDP-11/45 computer, which also serves as our principal speech processing facility. Recent work has focused on bringing up the new PDP-11/44 machine and moving all miniconcentrator activities to that facility. The PDP-11/44 miniconcentrator is now operating satisfactorily, and has been tested for full-duplex PCM speech traffic in a PVT-gateway-PVT communications mode. The LH/DH board*

* The LH/DH board can support an interface to the PSAT in either local host (LH) or distant host (DH) mode. Only the DH mode, which includes line drivers to and from the PSAT, has been used at Lincoln and ISI.

and PSAT connection have also been moved to the 11/44, and are operating successfully in that environment.

Some problems have been encountered with the use of the SIO chip on the UMC-Z80 for communication between the miniconcentrator and LEXNET. Occasional packets are damaged by what we suspect may be data overrun/underrun conditions that are partially detected by the SIO chip, but are not handled according to our understanding of the chip specifications. A new Z80 personality module for our logic analyzer is being used to investigate the problem in more detail. In the interim, software modifications have been made that allow error-free operation at reduced data rates.

As a result of these investigations, some changes have been made in the connections between the LEXNET interface and the UMC-Z80 in the miniconcentrator. The handshaking has been changed to facilitate full-duplex operation of the interface, and an additional wire has been added to allow the UMC-Z80 to send a RESET signal to the 8085 processors in the interface when starting operations. This change has allowed full-duplex operation at 375 kbps, but the UMC-Z80 specifications indicate that operation at at least twice this rate ought to be achievable. While some progress has been made, and our understanding of the operation of the SIO and associated DMA chips has grown, satisfactory performance at full rate has not yet been achieved. Current problems appear to be due to lost interrupts in the handshaking procedure. Work on the problem will continue into the next quarterly period.

It should be emphasized that current performance of the LEXNET-to-UMC-Z80 interface is sufficient to accommodate speech traffic from a number of PCM and LPC terminals, and to allow substantial system-level checkout of LEXNETs, gateways, and the wideband SATNET.

During this reporting period, a total of four additional protohex DH interface boards were fabricated, and all have been successfully debugged to integrate the Z80 complex with the PSAT. One of these cards was provided to ISI for use in their Z80 PSAT test bed, and a TELNET connection was established and used to debug the LH/DH card from the Lincoln facility.

C. Miniconcentrator Software

The various components of the gateway program relating to point-to-point connections became operational during this period. These components include an incoming message sorter, an ST message forwarder and aggregator, an IP forwarder, a connection manager, a message dispatcher, a LEXNET interface package, and a UMC-Z80 function simulator for the ARPANET. In its initial internet demonstration, the gateway was successfully used in an experiment in which ISI established a connection between a speech terminal and a speech echoer (both at ISI) through the gateway program running at Lincoln. ISI then sent speech (via ARPANET) to the gateway which forwarded the packets to the echoer. Runs were done both with an echoer which returned the speech by forwarding it back over the ARPANET and through the gateway a second time, and with an echoer which eliminated the second ARPANET loop in order to reduce throughput demands. This experiment was demonstrated at the November Wideband meeting with about 45 min. of continuous speech being transmitted through the gateway.

The miniconcentrator was successfully demonstrated at the end of January as a LEXNET/ARPANET gateway. Conversational speech using LPC encoding was sent from the LEXNET terminal through the UMC-Z80 to the PDP-11/45 where the gateway program forwarded the speech via the ARPANET to the terminal at ISI. The reverse path was exercised for speech initiating at ISI. Speech packets were forwarded through the ST message forwarder in the gateway,

and protocol packets were forwarded in IP form. Full ST point-to-point connection setup was carried out. This capability was demonstrated between Lincoln and ISI during the January ARPA Internet meeting held at ISI.

As mentioned above, the miniconcentrator project has moved onto the newly arrived PDP-11/44 from the PDP-11/45 which was used temporarily. A downloader program has been implemented for communicating via a terminal line with the PDP-11/44 from the PDP-11/70 which is serving as the development computer. The downloader permits one to type to the PDP-11/44 and receive responses. Also, it provides a mechanism for transferring files in a variety of formats from the PDP-11/70 to the PDP-11/44. In particular, the EPOS file system can be downloaded and started remotely from the PDP-11/70.

A terminal multiplexing capability has been incorporated in the EPOS operating system running on the PDP-11/44 and the "downloader" program running on the PDP-11/70. The resulting capability permits a user at the host computer (in this case the Lincoln PDP-11/70) to control up to six terminals in an interleaved fashion on the object miniconcentrator computer (in this case the Lincoln PDP-11/44). This capability will permit us to control and run experiments using miniconcentrators at sites other than Lincoln.

A program was written to download a LEXNET voice terminal via a terminal line from EPOS running on a PDP-11 computer. Heretofore, this operation could be performed only from the UNIX environment. Since ISI regularly runs EPOS rather than UNIX on their PDP-11s, the change was necessary to allow routine downloading of our software into the PVTs located at ISI. This program was utilized in downloading the ISI PVTs for their successful demonstration in March.

An EPOS user program was written which runs a fixed job in the background utilizing CPU time not used by other EPOS programs. Running another job (i.e., the gateway) in the foreground in competition with this timer program expands the running time of the timer's job and thereby provides a measure of the CPU usage of the foreground program. Preliminary runs with this timer on the Lincoln PDP-11/44 show that a quiescent gateway uses about 6 percent of the CPU; another 14 percent is used to support a PCM-encoded conversation (100 packets per second passing through the gateway).

Several changes have been made to the UMC-Z80 software that communicates with the LEXNET. Frequent receiver overruns in the UMC-Z80 were eliminated by modifying the handshaking between the UMC-Z80 and the LEXNET. The end-of-message handling scheme has been made simpler and faster. Some hardware problems referred to above have thus far prevented us from running the SIO chips at the 880-kbps full-duplex rate advertised by ACC, although the currently operable rate of 375 kbps is sufficient to support ongoing experiments.

The PSAT UMC-Z80 software has been rewritten in the same format as the LEXNET UMC-Z80 software. The original software had evolved from a series of progressively more complicated test programs for the 1822 interface. Converting the existing code to the LEXNET UMC-Z80 format enables us to use much of the same software in the PSAT UMC-Z80. In particular, the PDP-11 I/O developed for the LEXNET can be used on the PSAT side with little change. The PSAT UMC-Z80 modules which interact with the LH/DH interface and the PSAT have been debugged to the point where a periodic sequence of datagrams can be looped through the PSAT. The datagrams are produced by a datagram generator and receiver which simulates the presence of the PDP-11. Experiments to use this capability to determine the throughput of the combination of Z80, LH/DH interface, and PSAT host access module are now being initiated.

The next step will be to incorporate the software to communicate with the PDP-11 and to carry out a system-level checkout.

The PSAT UMC-Z80 software is interrupt driven on both the receive and transmit sides, using two Z80-DMA channels. System clock routines control the sending of miniconcentrator status information to the PSAT, and monitor the receipt of network status messages from the PSAT. Included in the software is a restart routine which performs the link initialization functions required by the PSAT Host Access Protocol (HAP), and an I/O initialization routine which sets up the DMA channels on the UMC-Z80 board and a Z80-PIO chip on the LH/DH board which is configured to function as a "Command and Status Register" between the UMC-Z80 and the LH/DH interface. The restart routine has recently been revised and updated based on ongoing interaction with BBN.

In the course of initial debugging, a receiver problem was traced to a need for a "PSAT transmitter inhibit" signal which required a slight design change to the LH/DH hardware. This change has been implemented on one of the LH/DH boards, and software tests indicate that it has corrected the problem.

D. Multiport Memory Design

The design of the Multiport Buffer Memory has been examined during this period in the light of expanding system requirements. The scope of the design has undergone a four-fold increase from the original 32K × 8 sizing to 64K × 16. Although this main memory expansion is significant in itself with regard to the impact on the hardware, the preliminary design study has resulted in a completely new set of architectural considerations. Whereas the original design concept was simply to supply a shared mass memory resource, the detailed system considerations have led to the requirement for additional intelligence within that resource. This need has arisen from an identification of the system bottlenecks in limiting the throughput of packet handling in the miniconcentrator complex.

Approximately 75 percent of the delay in packet handling in the miniconcentrator system environment relates to addressing overhead in the UMC-Z80. As a result of this realization, the architectural design of the Multiport Buffer Memory acting as a shared resource amongst a complex of four UMC-Z80 subsystems has been modified to incorporate the address manipulation necessary to unburden these complexes of that task. The design incorporates an intelligent memory management strategy in which packets of data destined for storage in the common resource are arranged in stringed blocks of 16 words. Each block is composed of 15 data words and a pointer to the address of the next data block in that packet. This organization facilitates the overlap of bus grant arbitration to multiple subsystems requesting simultaneous access to the shared memory. The general scenario for writing a new packet from any UMC-Z80 subsystem requesting access to the shared memory consists of sending the block length to the memory and receiving from it a corresponding address pointer to the beginning of that block. The UMC-Z80 then attaches the pointer to the header of the block before passing it on to the PDP-11/44 for system traffic control. The header then carries with it the location of the first pointer address to the string without having to retain any knowledge of how that string is distributed within the Multiport Memory. In this manner, strings which are destroyed after reading can be added to a "free list" of strings which may then be reassigned to fresh packets waiting to be written.

The major impact of the increased sizing of the memory has been the selection of a 65K dynamic RAM (MCM6664) in place of the faster 1K × 8 static RAM (MK4801) previously selected. The use of a 16-bit memory word allows the slower dynamic RAM to be used without adversely affecting the system throughput. The design will accommodate a worst case condition where all four UMC-Z80 complexes may simultaneously petition for access without holding off any single request beyond normal memory access latency. Data transfers between the Z80 and the memory are accomplished using byte-at-a-time DMA. A separate DMA for each complex as well as a separate PIO for control and status reporting is provided within the Multiport Memory. The full memory system will be accommodated on two protohex wirewrap boards.

V. EXPERIMENT DEFINITION AND PLANNING

The second annual supplement to the experiment plan was published as Lincoln Laboratory Project Report EWN-1, Supplement 2. It is also indexed as W-Note-25 in the Wideband Working Note series. The supplement provides updated information in the three subject areas of the experiment plan, namely network development, system validation, and advanced systems experiments. Some of the key events in system development and integration, as well as some current unresolved issues, are summarized briefly below.

There are now three nodes – Lincoln, ISI, and DCEC – which are fully equipped with PSAT, ESI, and Earth Station. The initial satellite system integration took place at ISI. Through a concerted effort by Western Union, Linkabit, and BBN, datagram loopback from the PSAT at ISI to the satellite and back was achieved during the 6-7 November Wideband Meeting. More details are given in the Wideband Meeting Summary which was prepared by Lincoln and distributed to the meeting attendees as W-Note-24. At that meeting, Lincoln and ISI demonstrated an important milestone in the effort to bring up concentrators and packet speech facilities, by linking LPC voice from ISI to Lincoln and back over the ARPANET, as described in Section IV.

In December, Linkabit delivered an ESI to Lincoln and proceeded (together with Western Union and BBN) to integrate and check out the PSAT, ESI, and Earth Station. Successful satellite loop tests were done under modem control, and the PSAT/ESI interface was almost completely checked out; the remaining WB SATNET validation is still in progress, but has been slowed by some system problems described below. A current equipment configuration at the Lincoln site is shown in Fig. 7. The illustrated configuration includes two PVTs, with an additional PVT used in traffic emulation mode. As discussed in Section IV-E, several more PVTs are under construction. A major milestone experiment, described in Section IV, was the LEXNET-to-ISI LPC communication over the ARPANET achieved in January. The current near-term goal is to achieve integration of the WB SATNET (i.e., reliable PSAT-to-PSAT communication) followed by packet speech communication between the Lincoln and ISI LEXNETs.

Two system problem areas have slowed the validation of the WB SATNET, with the result that the system is (as of 31 March 1981) not yet available for routine use. These problem areas are difficulty in establishing and demonstrating sustained performance at the specified channel bit error rate of 5×10^{-3} and persistent bugs in the remote fault monitoring and control equipment. The BER problem has several suspected causes, including insufficient transmitted power, unexpected losses in the receive chain, and intermodulation interference from other users on the same satellite transponder. The Linkabit burst test modems ordered by Western Union will help the situation, because they are more sensitive than the Scientific-Atlanta test modems presently in use; delivery of the first of the new modems is not expected, however, until mid-April.

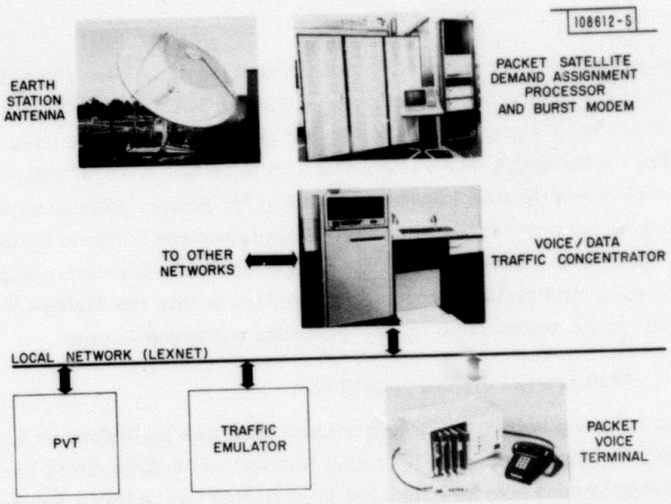

Fig. 7. Current wideband network experiment configuration at Lincoln site.

Meanwhile, a concerted effort is in progress by Western Union and site personnel to peak up the performance of the WB SATNET equipment and to carry out 48-hour continuous BER measurements, using S-A modems, for an appropriate set of loopback and site-to-site configurations.

The remote fault monitoring and control equipment is necessary in order to permit unmanned operation of the earth stations, by allowing Western Union's central satellite monitoring station in Glenwood, New Jersey, to detect and respond to any transmitter misbehavior (such as frequency drift) that violates FCC regulations and threatens other satellite users. The equipment includes fault sensors and automatic telephone dialing and answering equipment at the earth stations and at Glenwood. A minicomputer and interactive operator display facility at Glenwood can send touch-tone-coded command signals to the sites to shut down their transmitters. A series of problems has been encountered and resolved with this equipment, and full operation appears imminent. Up to the present time, tests over the satellite channel have been possible only when advance arrangements have been made for an FCC-licensed Western Union technician to be present at each earth station involved. Such arrangements have been made on several occasions to support PSAT and ESI tests; unfortunately, while the interfaces between these items appear to be functioning correctly, the tests have been largely inconclusive because of the satellite channel performance problems.

Lincoln has been working closely with Western Union, COMSAT, and the other contractors, in order to help expedite resolution of these problems. A note detailing current performance limitations and recommended courses of action was sent to DARPA in late March. Current indications are that channel availability on a routine basis is imminent.

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER: ESD-TR-81-87 — 2. GOVT ACCESSION NO.: AD-A104373	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): Packet Speech Systems Technology	5. TYPE OF REPORT & PERIOD COVERED: Semiannual Technical Summary, 1 October 1980 – 31 March 1981
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): Clifford J. Weinstein and Peter E. Blankenship	8. CONTRACT OR GRANT NUMBER(s): F19628-80-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Lincoln Laboratory, M.I.T., P.O. Box 73, Lexington, MA 02173	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: ARPA Order 3673, Program Element Nos. 62708E/61101E, Project Nos. 1D10/1T10
11. CONTROLLING OFFICE NAME AND ADDRESS: Defense Advanced Research Projects Agency, 1400 Wilson Boulevard, Arlington, VA 22209	12. REPORT DATE: 31 March 1981
	13. NUMBER OF PAGES: 28
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Electronic Systems Division, Hanscom AFB, Bedford, MA 01731	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

packet speech
network speech
voice conferencing
ARPANET
SATNET
LEXNET
digital channel vocoder

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes work performed on the Packet Speech Systems Technology Program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 October 1980 through 31 March 1981.